

Generic rigidity of reflection frameworks

Justin Malestein* Louis Theran†

Abstract

We give a combinatorial characterization of generic minimally rigid reflection frameworks. The main new idea is to study a pair of direction networks on the same graph such that one admits faithful realizations and the other has only collapsed realizations. In terms of infinitesimal rigidity, realizations of the former produce a framework and the latter certifies that this framework is infinitesimally rigid.

1. Introduction

A *reflection framework* is a planar structure made of *fixed-length bars* connected by *universal joints* with full rotational freedom. Additionally, the bars and joints are symmetric with respect to a reflection through a fixed axis. The allowed motions preserve the *length* and *connectivity* of the bars and *symmetry* with respect to some reflection. This model is very similar to that of *cone frameworks* that we introduced in [7]; the difference is that the symmetry group $\mathbb{Z}/2\mathbb{Z}$ acts on the plane by reflection instead of rotation through angle π .

When all the allowed motions are Euclidean isometries, a reflection framework is *rigid* and otherwise it is *flexible*. In this paper, we give a *combinatorial* characterization of minimally rigid, generic reflection frameworks.

1.1. The algebraic setup and combinatorial model Formally a reflection framework is given by a triple $(\tilde{G}, \varphi, \tilde{\ell})$, where \tilde{G} is a finite graph, φ is a $\mathbb{Z}/2\mathbb{Z}$ -action on \tilde{G} that is free on the vertices and edges, and $\tilde{\ell} = (\ell_{ij})_{ij \in E(\tilde{G})}$ is a vector of non-negative *edge lengths* assigned to the edges of \tilde{G} . A *realization* $\tilde{G}(\mathbf{p}, \Phi)$ is an assignment of points $\mathbf{p} = (\mathbf{p}_i)_{i \in V(\tilde{G})}$ and a representation of $\mathbb{Z}/2\mathbb{Z}$ by a reflection $\Phi \in \text{Euc}(2)$ such that:

$$\|\mathbf{p}_j - \mathbf{p}_i\|^2 = \ell_{ij}^2 \quad \text{for all edges } ij \in E(\tilde{G}) \quad (1)$$

$$\mathbf{p}_{\varphi(\gamma)\cdot i} = \Phi(\gamma) \cdot \mathbf{p}_i \quad \text{for all } \gamma \in \mathbb{Z}/2\mathbb{Z} \text{ and } i \in V(\tilde{G}) \quad (2)$$

The set of all realizations is defined to be the *realization space* $\mathcal{R}(\tilde{G}, \varphi, \ell)$ and its quotient by the Euclidean isometries $\mathcal{C}(\tilde{G}, \varphi, \ell) = \mathcal{R}(\tilde{G}, \varphi, \ell)/\text{Euc}(2)$ to be the configuration space. A realization is *rigid* if it is isolated in the configuration space and otherwise *flexible*.

As the combinatorial model for reflection frameworks it will be more convenient to use colored graphs. A *colored graph* (G, γ) is a finite, directed¹ graph G , with an assignment $\gamma =$

*Temple University, justmale@temple.edu

†Institut für Mathematik, Diskrete Geometrie, Freie Universität Berlin, theran@math.fu-berlin.de

¹For the group $\mathbb{Z}/2\mathbb{Z}$, the orientation of the edges do not play a role, but we give the standard definition for consistency.

$(\gamma_{ij})_{ij \in E(G)}$ of an element of a group Γ to each edge. In this paper Γ is always $\mathbb{Z}/2\mathbb{Z}$. There is a standard dictionary [7, Section 9] associating (\tilde{G}, φ) with a colored graph (G, γ) : G is the quotient of \tilde{G} by Γ , and the colors encode the covering map via a natural map $\rho : \pi_1(G, b) \rightarrow \Gamma$. In this setting, the choice of base vertex does not matter, and indeed, we may define $\rho : H_1(G, \mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$ and obtain the same theory.

1.2. Main Theorem We can now state the main result of this paper.

Theorem 1. *A generic reflection framework is minimally rigid if and only if its associated colored graph is reflection-Laman.*

The *reflection-Laman graphs* appearing in the statement are defined in Section 2. Genericity has its standard meaning from algebraic geometry: the set of non-generic reflection frameworks is a measure-zero algebraic set, and a small geometric perturbation of a non-generic reflection framework yields a generic one.

1.3. Infinitesimal rigidity and direction networks As in all known proofs of “Maxwell-Laman-type” theorems such as Theorem 1, we give a combinatorial characterization of a linearization of the problem known as *infinitesimal rigidity*. To do this, we use a *direction network* method (cf. [6, 7, 9, 10]). A *reflection direction network* $(\tilde{G}, \varphi, \mathbf{d})$ is a symmetric graph, along with an assignment of a *direction* \mathbf{d}_{ij} to each edge. The *realization space* of a direction network is the set of solutions $\tilde{G}(\mathbf{p})$ to the system of equations:

$$\langle \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij}^\perp \rangle = 0 \quad \text{for all edges } ij \in E(\tilde{G}) \quad (3)$$

$$\mathbf{p}_{\varphi(\gamma)\cdot i} = \Phi(\gamma) \cdot \mathbf{p}_i \quad \text{for all } \gamma \in \mathbb{Z}/2\mathbb{Z} \text{ and } i \in V(\tilde{G}) \quad (4)$$

where the $\mathbb{Z}/2\mathbb{Z}$ -action Φ on the plane is by reflection through the y -axis. A reflection direction network is determined by assigning a direction to each edge of the colored quotient graph (G, γ) of (\tilde{G}, φ) (cf. [7, Lemma 17.2]). Since all the direction networks in this paper are reflection direction networks, we will refer to them simply as “direction networks” to keep the terminology manageable. A realization of a direction network is *faithful* if none of the edges of its graph have coincident endpoints and *collapsed* if all the endpoints are coincident.

A basic fact in the theory of finite planar frameworks [3, 9, 10] is that, if a direction network has faithful realizations, the dimension of the realization space is equal to that of the space of infinitesimal motions of a generic framework with the same underlying graph. In [6, 7], we adapted this idea to the symmetric case when all the symmetries act by rotations and translations.

As discussed in [7, Section 1.8], this so-called “parallel redrawing trick”² described above does *not* apply verbatim to reflection frameworks. Thus, we rely on the somewhat technical (cf. [6, Theorem B], [7, Theorem 2]) Theorem 2, which we state after giving an important definition.

Let $(\tilde{G}, \varphi, \mathbf{d})$ be a direction network and define $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ to be the direction network with $(\mathbf{d}^\perp)_{ij} = (\mathbf{d}_{ij})^\perp$. These two direction networks form a *special pair* if:

- $(\tilde{G}, \varphi, \mathbf{d})$ has a faithful realization.
- $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ has only collapsed realizations.

²This terminology comes from the engineering community, in which the basic idea has been folklore for quite some time.

Theorem 2. Let (G, γ) be a colored graph with n vertices, $2n - 1$ edges, and lift (\tilde{G}, φ) . Then there are directions \mathbf{d} such that the direction networks $(\tilde{G}, \varphi, \mathbf{d})$ and $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ are a special pair if and only if (G, γ) is reflection-Laman.

Briefly, we will use Theorem 2 as follows: the faithful realization of $(\tilde{G}, \varphi, \mathbf{d})$ gives a symmetric immersion of the graph \tilde{G} that can be interpreted as a framework, and the fact that $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ has only collapsed realizations will imply that the only symmetric infinitesimal motions of this framework correspond to translation parallel to the reflection axis.

1.4. Notations and terminology In this paper, all graphs $G = (V, E)$ may be multi-graphs. Typically, the number of vertices, edges, and connected components are denoted by n , m , and c , respectively. The notation for a colored graph is (G, γ) , and a symmetric graph with a free $\mathbb{Z}/2\mathbb{Z}$ -action is denoted by (\tilde{G}, φ) . If (\tilde{G}, φ) is the lift of (G, γ) , we denote the fiber over a vertex $i \in V(G)$ by \tilde{i}_γ , with $\gamma \in \mathbb{Z}/2\mathbb{Z}$, and the fiber over a directed edge ij with color γ_{ij} by $\tilde{i}_\gamma \tilde{j}_{\gamma + \gamma_{ij}}$.

We also use (k, ℓ) -sparse graphs [5] and their generalizations. For a graph G , a (k, ℓ) -basis is a maximal (k, ℓ) -sparse subgraph; a (k, ℓ) -circuit is an edge-wise minimal subgraph that is not (k, ℓ) -sparse; and a (k, ℓ) -component is a maximal subgraph that has a spanning (k, ℓ) -graph.

Points in \mathbb{R}^2 are denoted by $\mathbf{p}_i = (x_i, y_i)$, indexed sets of points by $\mathbf{p} = (\mathbf{p}_i)$, and direction vectors by \mathbf{d} and \mathbf{v} . Realizations of a reflection direction network $(\tilde{G}, \varphi, \mathbf{d})$ are written as $\tilde{G}(\mathbf{p})$, as are realizations of abstract reflection frameworks. Context will always make clear the type of realization under consideration.

1.5. Acknowledgements LT is supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 247029-SDModels. JM is supported by NSF CDI-I grant DMR 0835586.

2. Reflection-Laman graphs

In this short section we introduce the combinatorial families of sparse colored graphs we use.

2.1. The map ρ Let (G, γ) be a $\mathbb{Z}/2\mathbb{Z}$ -colored graph. Since all the colored graphs in this paper have $\mathbb{Z}/2\mathbb{Z}$ colors, from now on we make this assumption and write simply “colored graph”. We recall two key definitions from [7].

The map $\rho : H_1(G, \mathbb{Z}) \rightarrow \mathbb{Z}/2\mathbb{Z}$ is defined on cycles by adding up the colors on the edges. (The directions of the edges don’t matter for $\mathbb{Z}/2\mathbb{Z}$ colors. Similarly, neither does the traversal order.) As the notation suggests, ρ extends to a homomorphism from $H_1(G, \mathbb{Z})$ to $\mathbb{Z}/2\mathbb{Z}$, and it is well-defined even if G is not connected.

2.2. Reflection-Laman graphs Let (G, γ) be a colored graph with n vertices and m edges. We define (G, γ) to be a *reflection-Laman graph* if: the number of edges $m = 2n - 1$, and for all subgraphs G' , spanning n' vertices, m' edges, c' connected components with non-trivial ρ -image and c'_0 connected components with trivial ρ -image

$$m' \leq 2n' - c' - 3c'_0 \tag{5}$$

This definition is equivalent to that of *cone-Laman graphs* in [7, Section 15.4]. The underlying graph G of a reflection-Laman graph is a $(2, 1)$ -graph.

2.3. Ross graphs and circuits Another family we need is that of *Ross graphs* (see [2] for an explanation of the terminology). These are colored graphs with n vertices, $m = 2n - 2$ edges, satisfying the sparsity counts

$$m' \leq 2n' - 2c' - 3c'_0 \quad (6)$$

using the same notations as in (5). In particular, Ross graphs (G, γ) have as their underlying graph, a $(2, 2)$ -graph G , and are thus connected [5].

A *Ross-circuit*³ is a colored graph that becomes a Ross graph after removing *any* edge. The underlying graph G of a Ross-circuit (G, γ) is a $(2, 2)$ -circuit, and these are also known to be connected [5], so, in particular, a Ross-circuit has $c'_0 = 0$, and thus satisfies (5) on the whole graph. Since (5) is always at least (6), we see that every Ross-circuit is reflection-Laman.

Because reflection-Laman graphs are $(2, 1)$ -graphs and subgraphs that are $(2, 2)$ -sparse are, in addition, Ross-sparse, we get the following structural result.

Proposition 2.1 ([8, Proposition 5.1],[2, Lemma 11]). *Let (G, γ) be a reflection-Laman graph. Then each $(2, 2)$ -component of G contains at most one Ross-circuit, and in particular, the Ross-circuits in (G, γ) are vertex disjoint.*

2.4. Reflection-(2, 2) graphs The next family of graphs we work with is new. A colored graph (G, γ) is defined to be a *reflection-(2, 2)* graph, if it has n vertices, $m = 2n - 1$ edges, and satisfies the sparsity counts

$$m' \leq 2n' - c' - 2c'_0 \quad (7)$$

using the same notations as in (5).

The relationship between Ross graphs and reflection-(2, 2) graphs we will need is:

Proposition 2.2. *Let (G, γ) be a Ross-graph. Then for either*

- *an edge ij with any color where $i \neq j$*
- *or a self-loop ℓ at any vertex i colored by 1*

the graph $(G + ij, \gamma)$ or $(G + \ell, \gamma)$ is reflection-(2, 2).

Proof. Adding ij with any color to a Ross (G, γ) creates either a Ross-circuit, for which $c'_0 = 0$ or a Laman-circuit with trivial ρ -image. Both of these types of graph meet this count, and so the whole of $(G + ij, \gamma)$ does as well. \square

It is easy to see that every reflection-Laman graph is a reflection-(2, 2) graph. The converse is not true.

Proposition 2.3. *A colored graph (G, γ) is a reflection-Laman graph if and only if it is a reflection-(2, 2) graph and no subgraph with trivial ρ -image is a $(2, 2)$ -block.* \square

Let (G, γ) be a reflection-Laman graph, and let G_1, G_2, \dots, G_t be the Ross-circuits in (G, γ) . Define the *reduced graph* (G^*, γ) of (G, γ) to be the colored graph obtained by contracting each G_i , which is not already a single vertex with a self-loop (this is necessarily colored 1), into a new vertex v_i , removing any self-loops created in the process, and then adding a new self-loop with color 1 to each of the v_i . By Proposition 2.1 the reduced graph is well-defined.

³The matroid of Ross graphs has more circuits, but these are the ones we are interested in here. See Section 2.4.

Proposition 2.4. *Let (G, γ) be a reflection-Laman graph. Then its reduced graph is a reflection-(2, 2) graph.*

Proof. Let (G, γ) be a reflection-Laman graph with t Ross-circuits with vertex sets V_1, \dots, V_t . By Proposition 2.1, the V_i are all disjoint. Now select a Ross-basis (G', γ) of (G, γ) . The graph G' is also a (2, 2)-basis of G , with $2n - 1 - t$ edges, and each of the V_i spans a (2, 2)-block in G' . The (k, ℓ) -sparse graph Structure Theorem [5, Theorem 5] implies that contracting each of the V_i into a new vertex v_i and discarding any self-loops created, yields a (2, 2)-sparse graph G^+ on n^+ vertices and $2n^+ - 1 - t$ edges. It is then easy to check that adding a self-loop colored 1 at each of the v_i produces a colored graph satisfying the reflection-(2, 2) counts (7) with exactly $2n^+ - 1$ edges. Since this is the reduced graph, we are done. \square

2.5. Decomposition characterizations A *map-graph* is a graph with exactly one cycle per connected component. A *reflection-(1, 1)* graph is defined to be a colored graph (G, γ) where G , taken as an undirected graph, is a map-graph and the ρ -image of each connected component is non-trivial.

Lemma 2.5. *Let (G, γ) be a colored graph. Then (G, γ) is a reflection-(2, 2) graph if and only if it is the union of a spanning tree and a reflection-(1, 1) graph.*

Proof. By [7, Lemma 15.1], reflection-(1, 1) graphs are equivalent to graphs satisfying

$$m' \leq n' - c'_0 \tag{8}$$

for every subgraph G' . Thus, (7) is

$$m' \leq (n' - c'_0) + (n' - c' - c'_0) \tag{9}$$

The second term in (9) is well-known to be the rank function of the graphic matroid, and the Lemma follows from the Edmonds-Rota construction [4] and the Matroid Union Theorem. \square

In the next section, it will be convenient to use this slight refinement of Lemma 2.5.

Proposition 2.6. *Let (G, γ) be a reflection-(2, 2) graph. Then there is a coloring γ' of the edges of G such that:*

- The ρ -image of every subgraph in (G, γ') is the same as in (G, γ) .
- There is a decomposition of (G, γ') as in Lemma 2.5 in which the spanning tree has all edges colored by the identity.

Proof. It is shown in [6, Lemma 2.2] that ρ is determined by its image on a homology basis of G . Thus, we may start with an arbitrary decomposition of (G, γ) into a spanning tree T and a reflection-(1, 1) graph X , as provided by Lemma 2.5, and define γ' by coloring the edges of T with the identity and the edges of X with the ρ -image of their fundamental cycle in T in (G, γ) . \square

Proposition 2.6 has the following re-interpretation in terms of the symmetric lift (\tilde{G}, φ) :

Proposition 2.7. *Let (G, γ) be a reflection-(2, 2) graph. Then for a decomposition, as provided by Proposition 2.6, into a spanning tree T and a reflection-(1, 1) graph X :*

- Every edge $ij \in T$ lifts to the two edges i_0j_0 and i_1j_1 . (In other words, the vertex representatives in the lift all lie in a single connected component of the lift of T .)
- Each connected component of X lifts to a connected graph.

3. Special pairs of reflection direction networks

We recall, from the introduction, that for reflection direction networks, $\mathbb{Z}/2\mathbb{Z}$ acts on the plane by reflection through the y -axis, and in the rest of this section $\Phi(\gamma)$ refers to this action.

3.1. The colored realization system The system of equations (3)–(4) defining the realization space of a reflection direction network $(\tilde{G}, \varphi, \mathbf{d})$ is linear, and as such has a well-defined dimension. Let (G, γ) be the colored quotient graph of (\tilde{G}, φ) .

To be realizable at all, the directions on the edges in the fiber over $ij \in E(G)$ need to be reflections of each other. Thus, we see that the realization system is canonically identified with the solutions to the system:

$$\langle \Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij} \rangle = 0 \quad \text{for all edges } ij \in E(G) \quad (10)$$

From now on, we will implicitly switch between the two formalisms when it is convenient.

3.2. Genericity Let (G, γ) be a colored graph with m edges. A statement about direction networks $(\tilde{G}, \varphi, \mathbf{d})$ is *generic* if it holds on the complement of a proper algebraic subset of the possible direction assignments, which is canonically identified with \mathbb{R}^{2m} . Some facts about generic statements that we use frequently are:

- Almost all direction assignments are generic.
- If a set of directions is generic, then so are all sufficiently small perturbations of it.
- If two properties are generic, then their intersection is as well.
- The maximum rank of (10) is a generic property.

3.3. Direction networks on Ross graphs We first characterize the colored graphs for which generic direction networks have strongly faithful realizations. A realization is *strongly faithful* if no two vertices lie on top of each other. This is a stronger condition than simply being faithful which only requires that edges not be collapsed.

Proposition 3.1. *A generic direction network $(\tilde{G}, \varphi, \mathbf{d})$ has a unique, up to translation and scaling, strongly faithful realization if and only if its associated colored graph is a Ross graph.*

To prove Proposition 3.1 we expand upon the method from [7, Section 20.2], and use the following proposition.

Proposition 3.2. *Let (G, γ) be a reflection-(2, 2) graph. Then a generic direction network on the symmetric lift (\tilde{G}, φ) of (G, γ) has only collapsed realizations.*

Since the proof of Proposition 3.2 requires a detailed construction, we first show how it implies Proposition 3.1.

3.4. Proof that Proposition 3.2 implies Proposition 3.1 Let (G, γ) be a Ross graph, and assign directions \mathbf{d} to the edges of G such that, for any extension $(G + ij, \gamma)$ of (G, γ) to a reflection-(2, 2) graph as in Proposition 2.2, \mathbf{d} can be extended to a set of directions that is generic in the sense of Proposition 3.2. This is possible because there are a finite number of such extensions.

For this choice of \mathbf{d} , the realization space of the direction network $(\tilde{G}, \varphi, \mathbf{d})$ is 2-dimensional. Since solutions to (10) may be scaled or translated in the vertical direction, all solutions to $(\tilde{G}, \varphi, \mathbf{d})$ are related by scaling and translation. It then follows that a pair of vertices in the fibers over i and j are either distinct from each other in all non-zero solutions to (10) or always coincide. In the latter case, adding the edge ij with any direction does not change the dimension of the solution space, no matter what direction we assign to it. It then follows that the solution spaces of generic direction networks on $(\tilde{G}, \varphi, \mathbf{d})$ and $(\widetilde{G + ij}, \varphi, \mathbf{d})$ have the same dimension, which is a contradiction by Proposition 3.2. \square

3.5. Proof of Proposition 3.2 It is sufficient to construct a specific set of directions with this property. The rest of the proof gives such a construction and verifies that all the solutions are collapsed. Let (G, γ) be a reflection-(2, 2) graph.

Combinatorial decomposition We apply Proposition 2.6 to decompose (G, γ) into a spanning tree T with all colors the identity and a reflection-(1, 1) graph X . For now, we further assume that X is connected.

Assigning directions Let \mathbf{v} be a direction vector that is not horizontal or vertical. For each edge $ij \in T$, set $\mathbf{d}_{ij} = \mathbf{v}$. Assign all the edges of X the vertical direction. Denote by \mathbf{d} this assignment of directions.

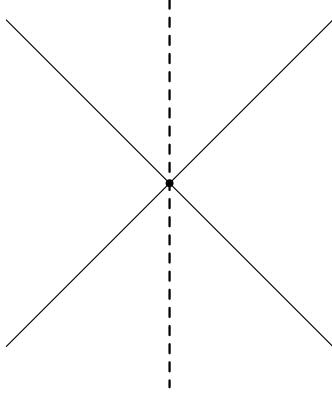


Figure 1: Schematic of the proof of Proposition 3.2: the y -axis is shown as a dashed line. The directions on the edges of the lift of the tree T force all the vertices to be on one of the two lines meeting at the y -axis, and the directions on the reflection-(1, 1) graph X force all the vertices to be on the y -axis.

All realizations are collapsed We now show that the only realizations of $(\tilde{G}, \varphi, \mathbf{d})$ have all vertices on top of each other. By Proposition 2.7 T lifts to two copies of itself, in \tilde{G} . It then

follows from the connectivity of T and the construction of \mathbf{d} that, in any realization, there is a line L with direction \mathbf{v} such that every vertex of \tilde{G} must lie on L or its reflection. Since the vertical direction is preserved by reflection, the connectivity of the lift of X , again from Proposition 2.7, implies that every vertex of \tilde{G} lies on a single vertical line, which must be the y -axis by reflective symmetry.

Thus, in any realization of $(\tilde{G}, \varphi, \mathbf{d})$ all the vertices lie at the intersection of L , the reflection of L through the y -axis and the y -axis itself. This is a single point, as desired. Figure 1 shows a schematic of this argument. \square

X does not need to be connected Finally, we can remove the assumption that X was connected by repeating the argument for each connected component of X separately. \square

3.6. Special pairs for Ross-circuits The full Theorem 2 will reduce to the case of a Ross-circuit.

Proposition 3.3. *Let (G, γ) be a Ross circuit with lift (\tilde{G}, φ) . Then there is an edge $i'j'$ such that, for a generic direction network $(\tilde{G}', \varphi, \mathbf{d}')$ with colored graph $(G - i'j', \gamma)$:*

- *The solution space of $(\tilde{G}', \varphi, \mathbf{d}')$ induces a well-defined direction \mathbf{d}_{ij} between i and j , yielding an assignment of directions \mathbf{d} to the edges of G .*
- *The direction networks $(\tilde{G}, \varphi, \mathbf{d})$ and $(\tilde{G}, \varphi, (\mathbf{d})^\perp)$ are a special pair.*

Before giving the proof, we describe the idea. We are after sets of directions that lead to faithful realizations of Ross-circuits. By Proposition 3.2, these directions must be non-generic. A natural way to obtain such a set of directions is to discard an edge ij from the colored quotient graph, apply Proposition 3.1 to obtain a generic set of directions \mathbf{d}' with a strongly faithful realization $\tilde{G}'(\mathbf{p})$, and then simply set the directions on the edges in the fiber over ij to be the difference vectors between the points.

Proposition 3.1 tells us that this procedure induces a well-defined direction for the edge ij , allowing us to extend \mathbf{d}' to \mathbf{d} in a controlled way. However, it does *not* tell us that rank of $(\tilde{G}, \varphi, \mathbf{d})$ will rise when the directions are turned by angle $\pi/2$, and this seems hard to do directly. Instead, we construct a set of directions \mathbf{d} so that $(\tilde{G}, \varphi, \mathbf{d})$ is rank deficient and has faithful realizations, and $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ is generic. Then we make a perturbation argument to show the existence of a special pair.

The construction we use is, essentially, the one used in the proof of Proposition 3.2 but turned through angle $\pi/2$. The key geometric insight is that horizontal edge directions are preserved by the reflection, so the “gadget” of a line and its reflection crossing on the y -axis, as in Figure 1, degenerates to just a single line.

3.7. Proof of Proposition 3.3 Let (G, γ) be a Ross-circuit; recall that this implies that (G, γ) is a reflection-Laman graph.

Combinatorial decomposition We decompose (G, γ) into a spanning tree T and a reflection-(1, 1) graph X as in Proposition 2.7. In particular, we again have all edges in T colored by the identity. For now, we *assume that X is connected*, and we fix $i'j'$ to be an edge that is on the cycle in X with $\gamma_{i'j'} \neq 0$; such an edge must exist by the hypothesis that X is reflection-(1, 1). Let $G' = G \setminus i'j'$. Furthermore, let T_0 and T_1 be the two connected components of the lift of T . For a

vertex $i \in G$, we denote the lift in T_0 by i_0 and the lift in T_1 by i_1 . We similarly denote the lifts of i' and j' by i'_0, i'_1 and j'_0, j'_1 .

Assigning directions The assignment of directions is as follows: to the edges of T , we assign a direction \mathbf{v} that is neither vertical nor horizontal. To the edges of X we assign the horizontal direction. Define the resulting direction network to be $(\tilde{G}, \varphi, \mathbf{d})$, and the direction network induced on the lift of G' to be $(\tilde{G}', \varphi, \mathbf{d})$.

The realization space of $(\tilde{G}, \varphi, \mathbf{d})$ Figure 2 contains a schematic picture of the arguments that follow.

Lemma 3.4. *The realization space of $(\tilde{G}, \varphi, \mathbf{d})$ is 2-dimensional and parameterized by exactly one representative in the fiber over the vertex i selected above.*

Proof. In a manner similar to the proof of Proposition 3.2, the directions on the edges of T force every vertex to lie either on a line L in the direction \mathbf{v} or its reflection. Since the lift of X is connected, we further conclude that all the vertices lie on a single horizontal line. Thus, all the points \mathbf{p}_{j_0} are at the intersection of the same horizontal line and L or its reflection. These determine the locations of the \mathbf{p}_{j_1} , so the realization space is parameterized by the location of $\mathbf{p}_{i'_0}$. \square

Inspecting the argument more closely, we find that:

Lemma 3.5. *In any realization $\tilde{G}(\mathbf{p})$ of $(\tilde{G}, \varphi, \mathbf{d})$, all the \mathbf{p}_{j_0} are equal and all the \mathbf{p}_{j_1} are equal.*

Proof. Because the colors on the edges of T are all zero, it lifts to two copies of itself, one of which spans the vertex set $\{\tilde{j}_0 : j \in V(G)\}$ and one which spans $\{\tilde{j}_1 : j \in V(G)\}$. It follows that in a realization, we have all the \mathbf{p}_{j_0} on L and the \mathbf{p}_{j_1} on the reflection of L . \square

In particular, because the color $\gamma_{i'j'}$ on the edge $i'j'$ is 1, we obtain the following.

Lemma 3.6. *The realization space of $(\tilde{G}, \varphi, \mathbf{d})$ contains points where the fiber over the edge $i'j'$ is not collapsed.*

The realization space of $(\tilde{G}', \varphi, \mathbf{d})$ The conclusion of Lemma 3.4 implies that the realization system for $(\tilde{G}, \varphi, \mathbf{d})$ is rank deficient by one. Next we show that removing the edge $i'j'$ results in a direction network that has full rank on the colored graph (G', γ) .

Lemma 3.7. *The realization space of $(\tilde{G}, \varphi, \mathbf{d})$ is canonically identified with that of $(\tilde{G}', \varphi, \mathbf{d})$.*

Proof. In the proof of Lemma 3.4, that X lifts to a connected subgraph of \tilde{G} was not essential. Because a horizontal line is preserved by the reflection, realizations will take on the same structure provided that X lifts to a subgraph with two connected components. Removing $i'j'$ from X leaves a graph X' with this property since X' is a tree.

It follows that the equation corresponding to the edge $i'j'$ in (10) was dependent. \square

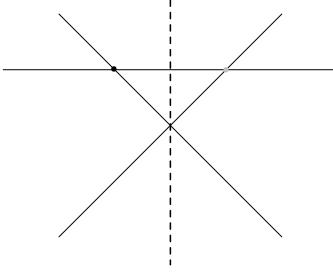


Figure 2: Schematic of the proof of Proposition 3.3: the y -axis is shown as a dashed line. The directions on the edges of the lift of the tree T force all the vertices to be on one of the two lines meeting at the y -axis. The horizontal directions on the connected reflection- $(1, 1)$ graph X force the point \mathbf{p}_{j_0} to be at the intersection marked by the black dot and \mathbf{p}_{j_1} to be at the intersection marked by the gray one.

The realization space of $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ Next, we consider what happens when we turn all the directions by $\pi/2$.

Lemma 3.8. *The realization space of $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ has only collapsed solutions.*

Proof. This is exactly the construction used to prove Proposition 3.2. \square

Perturbing $(\tilde{G}, \varphi, \mathbf{d})$ To summarize what we have shown so far:

- (a) $(\tilde{G}, \varphi, \mathbf{d})$ has a 2-dimensional realization space parameterized by $\mathbf{p}_{i'_0}$ and identified with that of a full-rank direction network on the Ross graph (G', γ) .
- (b) There are points $\tilde{G}(\mathbf{p})$ in this realization space where $\mathbf{p}_{i'_0} \neq \mathbf{p}_{j'_1}$.
- (c) $(\tilde{G}, \varphi, \mathbf{d})$ has a 1-dimensional realization space containing only collapsed solutions.

What we have not shown is that the realization space of $(\tilde{G}, \varphi, \mathbf{d})$ has *faithful* realizations, since the ones we constructed all have many coincident vertices. Proposition 3.1 will imply the rest of the theorem, provided that the above properties hold for any small perturbation of \mathbf{d} , since some small perturbation of *any* assignment of directions to the edges of (G', γ) has only faithful realizations.

Lemma 3.9. *Let $\hat{\mathbf{d}'}$ be a perturbation of the directions \mathbf{d}' on the edges of G' . If $\hat{\mathbf{d}'}$ is sufficiently close to \mathbf{d}' , then there are realizations of the direction network $(\tilde{G}', \varphi, \hat{\mathbf{d}'})$ such that $\mathbf{p}_{i'_0} \neq \mathbf{p}_{j'_1}$.*

Proof. The realization space is parameterized by $\mathbf{p}_{i'_0}$, and so $\mathbf{p}_{j'_1}$ varies continuously with the directions on the edges and $\mathbf{p}_{i'_0}$. Since there are realizations of $(\tilde{G}', \varphi, \mathbf{d})$ with $\mathbf{p}_{i_0} \neq \mathbf{p}_{j_1}$, the Lemma follows. \square

Lemma 3.9 implies that any sufficiently small perturbation of the directions assigned to the edges of G' gives a direction network that induces a well-defined direction on the edge $i'j'$ which is itself a small perturbation of $\mathbf{d}_{i'j'}$. Since the ranks of $(\tilde{G}', \varphi, \mathbf{d}')$ and $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ are stable under small perturbations, this implies that we can perturb \mathbf{d}' to a $\hat{\mathbf{d}'}$ that is generic in the sense of Proposition 3.1, while preserving faithful realizability of $(\tilde{G}, \varphi, \hat{\mathbf{d}'})$ and full rank of the realization system for $(\tilde{G}, \varphi, \hat{\mathbf{d}}^\perp)$. The Proposition is proved for when X is connected.

X need not be connected The proof is then complete once we remove the additional assumption that X was connected. Let X have connected components X_1, X_2, \dots, X_c . For each of the X_i , we can identify an edge $(i'j')_k$ with the same properties as $i'j'$ above.

Assign directions to the tree T as above. For X_1 , we assign directions exactly as above. For each of the X_k with $k \geq 2$, we assign the edges of $X_k \setminus (i'j')_k$ the horizontal direction and $(i'j')_k$ a direction that is a small perturbation of horizontal.

With this assignment \mathbf{d} we see that for any realization of $(\tilde{G}, \varphi, \mathbf{d})$, each of the X_k , for $k \geq 2$ is realized as completely collapsed to a single point at the intersection of the line L and the y -axis. Moreover, in the direction network on \mathbf{d}^\perp , the directions on these X_i are a small perturbation of the ones used on X in the proof of Proposition 3.2. From this follows that, in any realization $(\tilde{G}, \varphi, \mathbf{d}^\perp)$, is completely collapsed and hence full rank.

We now see that this new set of directions has properties (a), (b), and (c) above required for the perturbation argument. Since that argument makes no reference to the decomposition, it applies verbatim to the case where X is disconnected. \square

3.8. Proof of Theorem 2

The easier direction to check is necessity.

The Maxwell-direction If (G, γ) is not reflection-Laman, then it contains either a Laman-circuit with trivial ρ -image, or a violation of $(2, 1)$ -sparsity. If there is a Laman-circuit with trivial ρ -image, the Parallel Redrawing Theorem [11, Theorem 4.1.4] in the form [9, Theorem 3] implies that this subgraph has no faithful realizations for (G, φ, \mathbf{d}) only if it does in $(G, \varphi, \mathbf{d}^\perp)$ if rank-deficient. A violation of $(2, 1)$ -sparsity implies that the realization system (10) of $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ has a dependency, since the realization space is always at least 1-dimensional.

The Laman direction Now let (G, γ) be a reflection-Laman graph and let (G', γ) be a Ross-basis of (G, γ) . For any edge $ij \notin G'$, adding it to G' induces a Ross-circuit which contains some edge $i'j'$ having the property specified in Proposition 3.3. Note that $G' - ij + i'j'$ is again a Ross-basis. We therefore can assume (after edge-swapping in this manner) for all $ij \notin G'$ that ij has the property from Proposition 3.3 in the Ross-circuit it induces.

We assign directions \mathbf{d}' to the edges of G' such that:

- The directions on each of the intersections of the Ross-circuits with G' are generic in the sense of Proposition 3.3.
- The directions on the edges of G' that remain in the reduced graph (G^*, γ) are perpendicular to an assignment of directions on G^* that is generic in the sense of Proposition 3.2.
- The directions on the edges of G' are generic in the sense of Proposition 3.1.

This is possible because the set of disallowed directions is the union of a finite number of proper algebraic subsets in the space of direction assignments. Extend to directions \mathbf{d} on G by assigning directions to the remaining edges as specified by Proposition 3.3. By construction, we know that:

Lemma 3.10. *The direction network $(\tilde{G}, \varphi, \mathbf{d})$ has faithful realizations.*

Proof. The realization space is identified with that of $(\tilde{G}', \varphi, \mathbf{d}')$, and \mathbf{d}' is chosen so that Proposition 3.1 applies. \square

Lemma 3.11. *In any realization of $(\tilde{G}, \varphi, \mathbf{d}^\perp)$, the Ross-circuits are realized with all their vertices coincident and on the y -axis.*

Proof. This follows from how we chose \mathbf{d} and Proposition 3.3. \square

As a consequence of Lemma 3.11, and the fact that we picked \mathbf{d} so that \mathbf{d}^\perp extends to a generic assignment of directions $(\mathbf{d}^*)^\perp$ on the reduced graph (G^*, γ) we have:

Lemma 3.12. *The realization space of $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ is identified with that of $(\tilde{G}^*, \varphi, (\mathbf{d}^*)^\perp)$ which, furthermore, contains only collapsed solutions.*

Observe that a direction network for a single self-loop (colored 1) with a generic direction only has solutions where vertices are collapsed and on the y -axis. Consequently, replacing a Ross-circuit with a single vertex and a self-loop yields isomorphic realization spaces. Since the reduced graph is reflection-(2, 2) by Proposition 2.4 and the directions assigned to its edges were chosen generically for Proposition 3.2, that $(\tilde{G}, \varphi, \mathbf{d}^\perp)$ has only collapsed solutions follows. Thus, we have exhibited a special pair, completing the proof. \square

Remark It can be seen that the realization space of a direction network as supplied by Theorem 2 has at least one degree of freedom for each edge that is not in a Ross basis. Thus, the statement cannot be improved to, e.g., a unique realization up to translation and scale.

4. Infinitesimal rigidity of reflection frameworks

Let $(\tilde{G}, \varphi, \ell)$ be a reflection framework and let (G, γ) be the quotient graph. The configuration space, which is the set of solutions to the quadratic system (1)–(2) is canonically identified with the solutions to:

$$\|\Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i\|^2 = \ell_{ij}^2 \quad \text{for all edges } ij \in E(G) \quad (11)$$

where Φ acts on the plane by reflection through the y -axis. (That “pinning down” Φ does not affect the theory is straightforward from the definition of the configuration space: it simply removes rotation and translation in the x -direction from the set of trivial motions.)

4.1. Infinitesimal rigidity Computing the formal differential of (11), we obtain the system

$$\langle \Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i, \mathbf{v}_j - \mathbf{v}_i \rangle = 0 \quad \text{for all edges } ij \in E(G) \quad (12)$$

where the unknowns are the *velocity vectors* \mathbf{v}_i . A standard kind of result (cf. [1]) is the following.

Proposition 4.1. *Let $\tilde{G}(\mathbf{p}, \Phi)$ be a realization of an abstract framework $(\tilde{G}, \varphi, \ell)$. If the corank of the system (12) is one, then $\tilde{G}(\mathbf{p})$ is rigid.*

Thus, we define a realization to be *infinitesimally rigid* if the system (12) has maximal rank, and *minimally infinitesimally rigid* if it is infinitesimally rigid but ceases to be so after removing any edge from the colored quotient graph.

By definition, infinitesimal rigidity is defined by a polynomial condition in the coordinates of the points \mathbf{p}_i , so it is a generic property associated with the colored graph (G, γ) .

4.2. Relation to direction networks Here is the core of the direction network method for reflection frameworks: we can understand the rank of (12) in terms of a direction network.

Proposition 4.2. *Let $\tilde{G}(\mathbf{p}, \Phi)$ be a realization of a reflection framework with Φ acting by reflection through the y -axis. Define the direction \mathbf{d}_{ij} to be $\Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i$. Then the rank of (12) is equal to that of (10) for the direction network $(G, \gamma, \mathbf{d}^\perp)$.*

Proof. Exchange the roles of \mathbf{v}_i and \mathbf{p}_i in (12). □

4.3. Proof of Theorem 1 The, more difficult, “Laman direction” of the Main Theorem follows immediately from Theorem 2 and Proposition 4.2: given a reflection-Laman graph Theorem 2 produces a realization with no coincident endpoints and a certificate that (12) has corank one. □

4.4. Remarks The statement of Proposition 4.2 is *exactly the same* as the analogous statement for orientation-preserving cases of this theory. What is different is that, for reflection frameworks, the rank of $(G, \gamma, \mathbf{d}^\perp)$ is *not*, the same as that of (G, γ, \mathbf{d}) . By Proposition 3.2, the set of directions arising as the difference vectors from point sets are *always non-generic* on reflection-Laman graphs, so we are forced to introduce the notion of a special pair as in Section 3.

References

- [1] L. Asimow and B. Roth. The rigidity of graphs. *Trans. Amer. Math. Soc.*, 245:279–289, 1978. ISSN 0002-9947. doi: 10.2307/1998867. URL <http://dx.doi.org/10.2307/1998867>.
- [2] Matthew Berardi, Brent Heeringa, Justin Malestein, and Louis Theran. Rigid components in fixed-lattice and cone frameworks. In *Proceedings of the 23rd Annual Canadian Conference on Computational Geometry (CCCG)*, 2011. URL <http://arxiv.org/abs/1105.3234>.
- [3] Mike Devlin, Jeremy L. Martin, and Victor Reiner. Rigidity theory for matroids. *Comment. Math. Helv.*, 82(1):197–233, 2007. ISSN 0010-2571. doi: 10.4171/CMH/89. URL <http://dx.doi.org/10.4171/CMH/89>.
- [4] Jack Edmonds and Gian-Carlo Rota. Submodular set functions (abstract). In *Waterloo Combinatorics Conference*, University of Waterloo, Ontario, 1966.
- [5] Audrey Lee and Ileana Streinu. Pebble game algorithms and sparse graphs. *Discrete Math.*, 308(8):1425–1437, 2008. ISSN 0012-365X. doi: 10.1016/j.disc.2007.07.104. URL <http://dx.doi.org/10.1016/j.disc.2007.07.104>.
- [6] Justin Malestein and Louis Theran. Generic combinatorial rigidity of periodic frameworks. Preprint, arXiv:1008.1837, 2010. URL <http://arxiv.org/abs/1008.1837>.
- [7] Justin Malestein and Louis Theran. Generic rigidity of frameworks with orientation-preserving crystallographic symmetry. Preprint, arXiv:1108.2518, 2011. URL <http://arxiv.org/abs/1108.2518>.

- [8] Justin Malestein and Louis Theran. Generic rigidity with forced symmetry and sparse colored graphs. Preprint, arXiv:1203.0772, 2012. URL <http://arxiv.org/abs/1203.0772>.
- [9] Ileana Streinu and Louis Theran. Slider-pinning rigidity: a Maxwell-Laman-type theorem. *Discrete & Computational Geometry*, 44(4):812–837, 2010. ISSN 0179-5376. doi: 10.1007/s00454-010-9283-y. URL <http://dx.doi.org/10.1007/s00454-010-9283-y>.
- [10] Walter Whiteley. The union of matroids and the rigidity of frameworks. *SIAM J. Discrete Math.*, 1(2):237–255, 1988. ISSN 0895-4801. doi: 10.1137/0401025. URL <http://dx.doi.org/10.1137/0401025>.
- [11] Walter Whiteley. Some matroids from discrete applied geometry. In J. Bonin, James G. Oxley, and B. Servatius, editors, *Matroid Theory*, volume 197 of *Contemporary Mathematics*, pages 171–311. American Mathematical Society, 1996.